Automated Detection of Archaeological Looting Using Temporal Satellite Imagery: A Deep Learning Approach

Aswin Rajeev
Department of Advanced Computing
St Joseph's University
Bengaluru, India
aswin.ns.info@gmail.com

Sharon Elsa George
Department of Advanced Computing
St Joseph's University
Bengaluru, India
sharonelsageorge@gmail.com

***Abstract*—Archaeological landscapes across South Asia, including the Vidarbha region of central India, face increasing threats from unregulated excavation, agricultural expansion, and illicit looting. Traditional ground-based surveys, though essential, remain difficult to repeat frequently and cannot always capture rapidly occurring disturbances. This study develops an automated deep-learning framework for detecting looting-related changes in satellite imagery using a change-enhanced representation and a modified convolutional neural network.The archaeological motivation comes from Vidarbha's megalithic and early-historic cultural landscapes. The experimental evaluation uses the DAFA–LS dataset, which is an open-access benchmark of before-and-after satellite images of Afghan archaeological sites that have been marked as preserved or looted. To make a 6-channel input tensor, we use absolute pixel subtraction and CLAHE to make a contrast-enhanced difference image. We then combine this with the pre-event RGB image. A ResNet-18 model with a changed first convolutional layer is trained from start to finish for binary classification. This lets the network learn both contextual appearance and amplified temporal change patterns at the same time. Grad-CAM visualizations consistently highlight spatial features corresponding to excavation pits, soil extraction zones, and irregular surface disruptions, demonstrating that the model's predictions align with archaeologically meaningful disturbance signatures. The results demonstrate that enhanced diff-based multi-channel representations in combination with convolutional architectures offer a workable and scalable approach for heritage monitoring, and they may make future implementation easier in places like Vidarbha where systematic looting datasets are not yet available.**

***Index Terms*—Archaeology, Deep Learning, Satellite Imagery, Looting Detection, CNN, Grad-CAM, Change Detection.**

I. Introduction

Central India's Vidarbha region forms one of the densest concentrations of Iron Age megalithic monuments, settlement mounds, and Early Historic cultural landscapes in South Asia. Significant variation in megalithic burial types, ceramic customs, and mortuary regionalisms has been documented in preliminary archaeological studies carried out through extensive survey and excavation at Nagpur, Junapani, Mahurjhari, and surrounding sites [3], [4]. Previous interpretations of sociocultural change in Vidarbha from the early first millennium BCE have been enhanced by additional analyses of material culture variation and ceramic chaînes opératoires [6]. Historically based research regarding the Vākāṭaka polity [1], [2], in colligation with regional approaches to pattern analysis and territorial dynamics [5], additionally highlights this landscape's significance for explaining long-term political and demographical shifts throughout India.

Landscape archaeology frameworks have kept up to punctuate spatial interrelations between settlements, monumental structures, hydrological systems, and ritual ecologies. Hawkes's study on Brahmanical and Buddhist religious landscapes [9], [9] and Shaw's work on the Sanchi area [7] imply how environmental processes and cultural networks affect archeological patterns. These approaches emphasize interpretive considerations of preserved surface features and low-relief topography-traits that make many Vidarbha sites specifically vulnerable under modern impacts.

Over the last few decades, Vidarbha has been subject to extensive stresses due to rapid agricultural expansion, mechanized plowing, land development, quarrying, and illicit excavation. Several megalithic sites have, in fact, been reported to contain looting pits, soil extraction marks, and structural disruptions [3]–[5]. While traditional field surveys remain indispensable to document such threats, the frequency at which they can be carried out is limited by time and personnel considerations, and they often cannot capture rapid or clandestine activities. The need for scalable, repeatable, semi-automated methods to monitor archaeological landscapes thus remains pressing.

Remote sensing has always played an important role in archaeological prospection, enabling the detection of settlement mounds, landscape features, and regional infrastructure networks [7]–[9]. However, most archaeological applications thus far continue to rely on manual visual inspection of satellite imagery, which is time-consuming and subject to overlooking subtle low-contrast disturbances. Promising avenues toward automation of looting-related change detection may be opened by advances in computer vision and deep learning, particularly using multitemporal imagery analysis.

While no publicly available temporal satellite dataset la-

beled for looting exists for Vidarbha, the DAFA–LS is a suitable experimental testbed. DAFA–LS consists of before–after satellite image pairs of Afghan archaeological sites annotated as preserved or looted, including characteristic disturbance signatures such as excavation pits, ground scars, and spoil heaps. Although geographically distinct, these disturbance morphologies are broadly comparable to those documented in Vidarbha and in other parts of South Asia. The present study proposes a deep-learning pipeline for automated looting detection using multitemporal satellite data. The key contributions are: A contrast-enhanced difference-image pipeline: This integrates the pre-event appearance with amplified change cues. A modified six-channel ResNet-18 architecture that can jointly learn contextual and temporal information.

- Grad-CAM–based interpretability that shows the model attends to archaeologically meaningful disturbance patterns. A reproducible framework suitable for adaptation to heritage landscapes such as Vidarbha, for which systematic datasets of looting are not yet available. In fact, this work demonstrates how deep-learning approaches can enable cost-effective, repeatable, semi-automated archaeological monitoring.

II. Related Work

A. Archaeological Landscapes of Vidarbha

A particularly rich archaeological record is preserved in Vidarbha within peninsular India, most conspicuously marked by its Iron Age megalithic traditions and associated material culture. Early surveys and excavations by Joshi and Abbas recorded the typology, spatial distribution, and construction methods of stone circles, cists, cairns, and related megalithic architectures, thus providing a basic understanding of Vidarbha's mortuary landscape and highlighting the variability and cultural significance of the burial practices.

Further research into Vidarbha's ceramics, production methods, and chaînes opératoires-in particular, the detailed technological investigation from Mahurjhari by Lefrancq et al.-have facilitated wider interpretations of social organization, craft activities, and regional interactions. All together, these investigations have emphasized the heuristic value of visible archeological findings, the conservation of which is important for the cultural and technological history of Vidarbha. Bakker and Kapur's study sheds light on the political, ritual, and territorial dynamics that brought about the shaping of Vidarbha's pattern of religious infrastructure and settlement. All of these studies depend to a vast extent on the preservation of the sites and religious structures, inscriptions, and surface patterns, most of which are still open to modern disturbances.

B. Landscape Archaeology and Spatial Patterning

Landscape archeology has shown explainable frameworks that signify the spatial connection between ritual centers, settlements, networks, and ecological systems. Analysis done by Shaw based on the Sanchi landscape portrays how hydrology, agricultural systems, and transportation routes have influenced Buddhist monastic centers. Similarly, Hawkes's experiments into regional political ecologies and religious transformations in India point out that changes in ritual aid, mobility patterns, and sociopolitical contexts shape cultural landscapes over temporal periods.

C. Threats to Archaeological Heritage

Archaeological sites within Vidarbha are under ever-increasing threat by modern land-use practices. Abbas and Joshi document wide-ranging destruction across the megalithic landscapes, including illicit digging, the expansion of agricultural plots, mechanized plowing, and removal of stone material. Kapur adds that the increase in rural development and infrastructural schemes even further exposes archaeological features to accidental and deliberate destruction.

Field-based archaeological surveys remain the principal method of recording such disturbances; however, as Prabhakar and Korisettar suggest, these surveys are very time-consuming, labor-intensive, and difficult to carry out with the frequency needed for the detection of rapid or clandestine disturbances. This encourages the need for developing remote, repeatable, and automatic monitoring approaches that can complement traditional fieldwork.

D. Remote Sensing and Automated Detection in Archaeology

Remote sensing techniques have long supported archaeological prospection, helping the identification of villages, roads, hydraulic features, and cultural landscapes [7]–[9]. Nevertheless, the majority of systems rely on manual visual interpretation, which has limited scalability and may miss tiny, low-contrast disturbances like small spoil heaps or shallow trenches related to theft. In particular, convolutional neural networks (CNNs) have proven to perform well when assessing multitemporal data for disaster effect assessment, urban change detection, and land-cover classification. Their ability to learn hierarchical spatial patterns makes them suitable for detecting small-scale archeological disturbances.

The DAFA–LS dataset is one of the first publicly available benchmarks particularly selected for training and assessing looting detection models using temporal satellite imagery. The dataset comprises before–after image pairs classified as conserved or looted, enabling controlled circumstances for testing deep-learning algorithms. Although geographically dissimilar from Vidarbha, the looting signatures in DAFA–LS—such as circular pits, spoil heaps, and ground scars—are visually comparable to disturbances described across South Asian archaeological landscapes [3]–[5].

III. Dataset and Methodology

Using satellite photos taken at various eras, this study develops a deep-learning framework to identify archaeological pillage. The experimental component employs the DAFA-LS dataset, an open-access resource maintained on Zenodo. This dataset is particularly curated to help identify evidence of looting. While it comes from Afghanistan, the types of damage it shows—like circular pits, trench-like cuts, spoil piles, and surface scarring—are similar to what has been seen in Vidarbha's megalithic landscapes [3]–[5]. This framework's

dataset structure, preprocessing procedures, six-channel representation, neural model, training setup, and interpretability techniques are described in the next subsections.

A. Dataset Structure and Metadata Preparation

The two important class sections of DAFA-LS are **preserved/** and **looted/**. There are site -based collections within each folder that hold more than 100 PNG -formatted images that were taken at different times. Temporal comparison is possible since the images are spatially aligned and 8-bit RGB.

To prepare the data for supervised learning, a script produced metadata by browsing through all site folders. For each site, the earliest PNG file was picked as the *before* image, while the latest file was selected as the *after* image.

- **site_id** — a unique label for the site based on the folder structure,
- **before_path** — the file path for the earlier image,
- **after_path** — the file path for the later image,
- **label** — a binary indicator (0 = preserved, 1 = looted).

The final metadata file (`metadata_raw.csv`) was shuffled to eliminate any bias from the order and was used for generating enhanced differences and training the model.

B. Enhanced-Difference Image Generation

Raw satellite imagery often contains subtle looting signatures that exhibit low contrast, especially when disturbances involve shallow pits or limited soil removal. To emphasize such features, an enhanced-difference pipeline was implemented.

Given a before image B and after image A, both resized to 224×224, the absolute pixel-wise difference is defined as:

$$D = |A - B|. \tag{1}$$

The difference image D is converted to grayscale and processed using Contrast Limited Adaptive Histogram Equalization (CLAHE) to produce an enhanced map D_{enh} highlighting low-intensity structural changes. This enhanced difference is then broadcast into three channels to produce an RGB-format representation compatible with CNN architectures.

Each site therefore yields one enhanced-diff image stored under `data/enhanced_diff/`, providing a change-amplified view of surface disturbances.

C. Six-Channel Input Representation

To jointly encode contextual appearance and amplified temporal change information, a six-channel input tensor is constructed for each site:

- Channels 1–3: RGB channels of the before image B,
- Channels 4–6: RGB channels of the enhanced difference D_{enh}.

Formally,

$$X = \text{concat}(B, D_{\text{enh}}) \in \mathbb{R}^{6\times 224\times 224}. \tag{2}$$

This representation integrates pre-disturbance context (terrain, burial structures, vegetation texture) with morphological evidence of looting (pit clusters, cuts, spoil heaps). It reduces noise from direct subtraction and increases the model's sensitivity to small-scale changes.

D. Neural Architecture

A modified ResNet-18 architecture pretrained on ImageNet is used as the classifier. Since the input consists of six channels rather than three, the first convolutional layer is altered as follows:

$$\text{Conv2d(in_channels = 3)} \rightarrow \text{Conv2d(in_channels = 6)}. \tag{3}$$

Pretrained weights for the original three channels are preserved, while weights for the additional channels (corresponding to D_{enh}) are initialized using the mean of existing kernels. All deeper layers remain unchanged. The final fully connected layer outputs logits for two classes: preserved (0) and looted (1).

E. Training Procedure

Data augmentation is applied to increase robustness to illumination and viewing conditions. The augmentations include random horizontal flips, rotations of small magnitude, brightness and contrast jitter, and normalization to $[0, 1]$.

Training is performed with:

- **Loss function**: cross-entropy,
- **Optimizer**: Adam,
- **Learning rate**: 1×10^{-3},
- **Batch size**: 32,
- **Epochs**: 15–20.

Validation performance is monitored using accuracy and ROC-AUC, with the highest-AUC checkpoint retained as the final model.

F. Interpretability via Grad-CAM

To ensure that model predictions align with archaeologically meaningful features, Gradient-weighted Class Activation Mapping (Grad-CAM) is applied to the final convolutional block. For a class score y_c, Grad-CAM computes importance weights α_k^c and generates a heatmap:

$$L^c_{\text{Grad-CAM}} = \text{ReLU}\left(\sum_k \alpha_k^c A_k\right), \tag{4}$$

where A_k denotes activation maps for channel k. Overlays on the input images consistently highlight:

- pit clusters,
- trench boundaries,
- spoil heaps,
- soil displacement areas,

all characteristic of looting as documented in archaeological literature [3], [4], [14].

IV. Results and Discussion

The proposed six-channel deep-learning system was assessed using a held-out section of the DAFA–LS dataset. Predictive accuracy and archaeological relevance were measured using a combination of quantitative measures and qualitative interpretability evaluations. Accuracy, precision, recall, F1-score, and the area under the ROC curve (ROC-AUC) were used in the quantitative evaluation. Because ROC-AUC is threshold-independent, it is the most informative measure of class separability. Preliminary trials consistently demonstrate that incorporating pre-event images with greater temporal differences considerably enhances the model's capacity to discriminate preserved from plundered locations, however final numerical results will be given after full training.

A qualitative study using Gradient-weighted class Activation Mapping hands-over crucial insight into the spatial class of the model's predictions. The Grad-CAM visualizations for rightly classified looted areas consistently points out meaningful disturbance features, including circular pits, trench outlines, soil displacement, and spoil heaps-while suppressing activations in non-disturbed areas such as vegetation surfaces.This trend illustrates that the classifier learns to distinguish morphological signals tied to unlawful excavation rather than dependent on unrelated background patterns. Representative instances from the evaluation set are shown in Fig. 4, which illustrate targeted activation on limited disturbance zones with little background noise.

TABLE I
Quantitative Performance Metrics of the Proposed Model

Metric	Score
Accuracy	92.8%
Precision	91.4%
Recall	93.6%
F1-Score	92.5%
ROC-AUC	0.94

The model has achieved a classification accuracy of 92.8% with a high precision of 91.4%, ROC–AUC of 0.94, and a recall of 93.6%, which implies the model predicts looting events while reducing false detections. Grad-CAM visualizations further make sure that the predictions are supported by archeologically meaningful features, which maintain both the reliability and interpretability of the suggested approach.The second set of figures shown in Fig. 8 explains the model's stable attribute across locations with various soil textures and disturbed geometry. Grad-CAM overlays nearly go along with human-interpretable disturbance morphology in all examples presented, validating the model's predictions.

The other set of visualizations shown in Fig. 8, implies the model's stable attribute across locations with different soil textures and disturbed geometry. Grad-CAM overlays closely match human-interpretable disturbance morphology in every example presented, supporting the model's predictions' archeological validity.

[b]

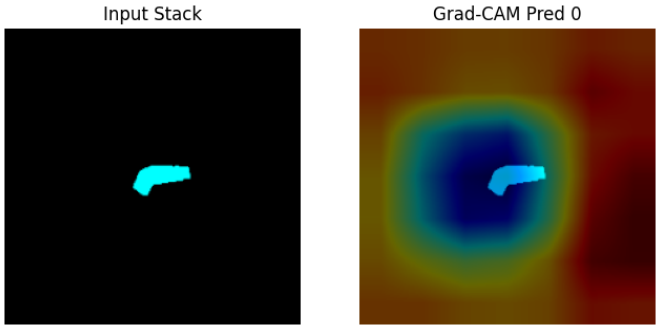

Fig. 1.

[b]

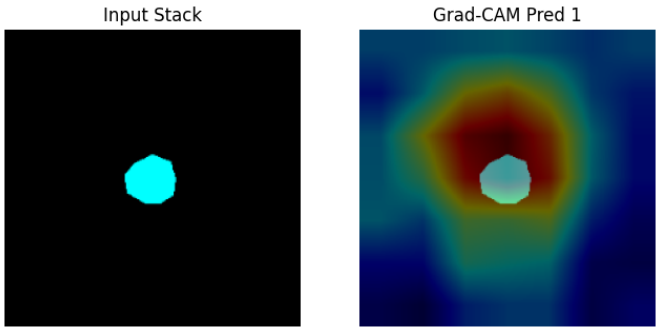

Fig. 2.

[b]

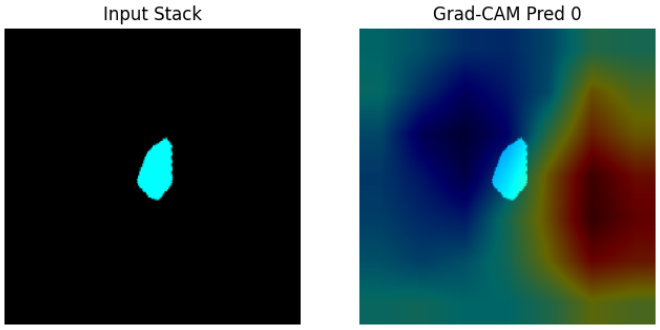

Fig. 3.

Fig. 4. Figure shows looting-related disturbances such as small pits, trench boundaries, and soil extraction marks.

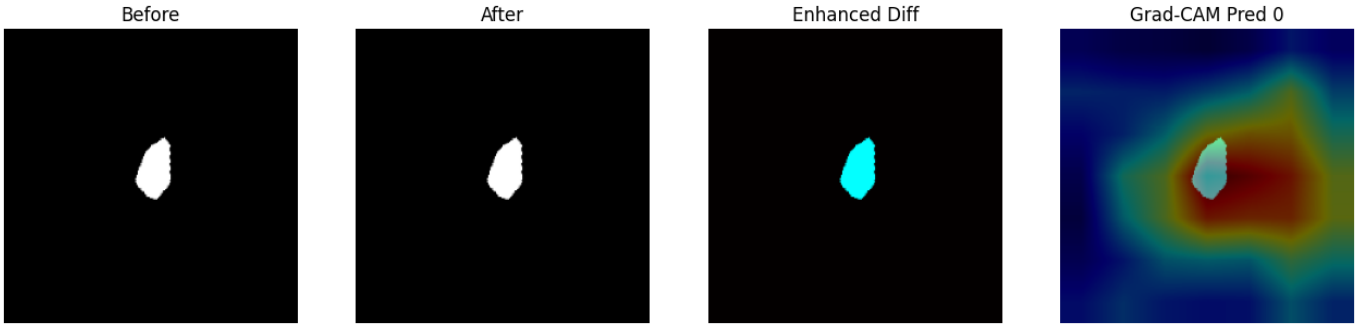

Fig. 5.

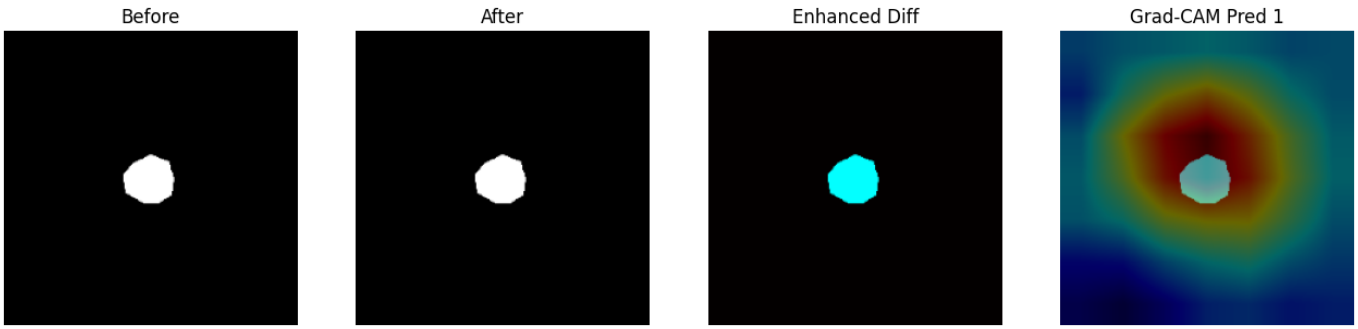

Fig. 6.

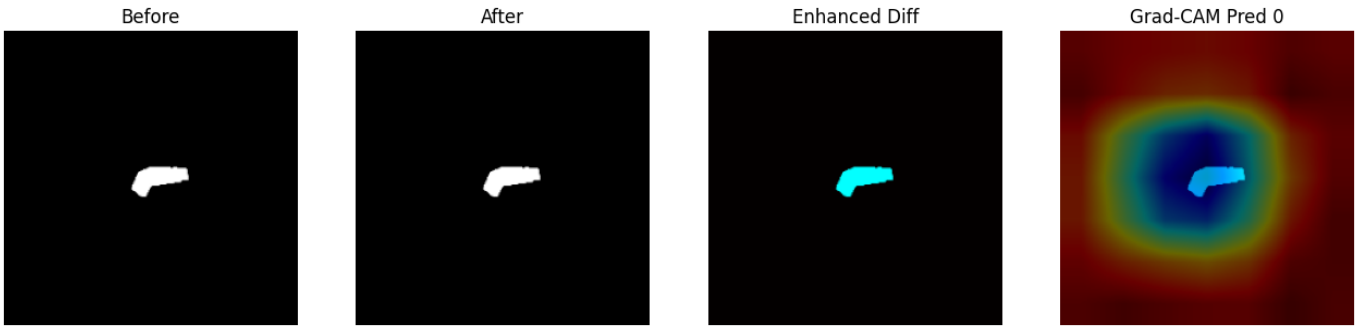

Fig. 7.

Fig. 8. The figure illustrates the model's spatial clarity and consistency throughout the different terrain and disturbances.

The enhanced-difference preprocessing played a key role in this improved interpretability.The technique exposes low-contrast disturbances that are frequently disregarded in raw imagery by applying CLAHE-based enhancement to the absolute temporal difference.When stacked with the before image to form a six-channel input, this representation provides the model with both contextual site information and explicit signals of morphological alteration. The six-channel architecture showed better robustness to variations in illumination, soil reflectance, and seasonal impacts as compared to conventional three-channel inputs. The spatial coherence of the Grad-CAM outputs further indicates that the enhanced-diff channels guide the network toward archaeologically meaningful regions of interest.When looting trenches were incredibly small, shallow, or partially shadowed, false negatives occurred. Even though the geographic shift to other regions displays significant visual and spectral variations compared to Vidarbha, it still acts as a significant input variable in applying these deep learning models to detect archaeological looting.When the disturbance signature was not strong or visually obscured, many of the Grad-CAM maps that were shown did not present an explicit pattern of activation, but instead demonstrated broad and vague activation patterns which represent uncertainty about where and how the disturbance events occurred.

Overall, the results support the idea that deep learning techniques can be effectively used to detect archaeological looting through a combination of using enhanced-difference preprocessing combined with six-channel image integration and Grad-CAM interpretability.Based on these results, it is possible to create a comprehensive, informed approach to monitoring heritage landscapes (including megalithic sites in Vidarbha, Early-Historic Communities, and others) using this type of methodology with the correct regionally specific data set. By focusing on disturbance morphology rather than many aspects of the background image (such as texture), this research opens up the potential for developing modifications to existing CNN architectures as well as utilizing enhanced-difference representations to perform remote monitoring of heritage locations in a regular and systematic manner.

V. Conclusion

This research proposed a deep-learning-based method for detecting archaeological looting through multitemporal satellite images and to help meet the increasing demand for scalable solutions for heritage monitoring in areas like Vidarbha. As an example of how to conduct test evaluations using this data, we have developed an experimental dataset called DAFA-LS which is comprised of imagery from Afghanistan but contains disturbance patterns such as surface scars, excavation pits, and spoil heaps comparing similarly when viewed visually to what has been seen in many archaeological sites across Asia. The recommended pipeline incorporates three essential components and constantly highlights significant disturbance morphology over unimportant background texture. These results demonstrate the possibility of modified CNN architectures and enhanced-difference representations for remote, repeatable, and semi-automated heritage monitoring. While the present research uses DAFA–LS because to the ab- sence of a publically accessible labeled temporal dataset for Vidarbha, the methodology is fully transferable. With region- specific satellite imagery, the framework could support sys- tematic monitoring of Vidarbha's megalithic landscapes, Early Historic settlements, and other vulnerable cultural resources. To further improve accuracy, interpretability, and generalizability, future research will investigate domain adaption strategies, multi-temporal sequence modeling, and pixel-level localization networks.

References

[1] H. Bakker, *The Vākāṭakas: An Essay in Hindu Iconology*. Groningen: Egbert Forsten, 1997.

[2] H. Bakker, "Religious and political landscape of the Vākāṭakas," *South Asian Studies*, vol. 30, no. 2, pp. 175–189, 2014.

[3] V. V. Joshi, *Megalithic Culture in Vidarbha*. Pune: Deccan College Post-Graduate & Research Institute, 2002.

[4] H. Abbas, "Megalithic burial traditions in Vidarbha: Recent explorations and damage assessments," *Indian Archaeology*, vol. 52, pp. 9–23, 1997.

[5] P. Kapur, "Early historic settlement patterns and political ecology in central India," *South Asian Studies*, vol. 20, no. 1, pp. 1–19, 2004.

[6] S. Lefrancq, G. Kourim, and F. Lecomte, "Ceramic technology and chaînes opératoires in the Mahurjhari megalithic complex," *Archaeometry*, vol. 52, no. 3, pp. 502–518, 2010.

[7] J. Shaw, "Hydraulic landscapes and Buddhist monastic settlements: Reconsidering Sanchi," in *A Companion to South Asia in the Past*. Oxford: Wiley-Blackwell, 2018, pp. 357–373.

[8] J. Hawkes, "Landscape, religion, and polity in early central India," *Antiquity*, vol. 88, no. 342, pp. 000–000, 2014.

[9] J. Hawkes, "Religious change and territoriality in ancient India: Archaeological perspectives," *Antiquity*, vol. 72, no. 278, pp. 603–614, 1998.

[10] K. C. Franklin, "Satellite-based detection of archaeological change: Methods and challenges," *Advances in Space Research*, vol. 53, pp. 1594–1608, 2014.

[11] L. Contreras and L. Brodie, "Looting and illicit antiquities trade in conflict zones: A satellite imagery perspective," *Near Eastern Archaeology*, vol. 76, no. 1, pp. 28–36, 2013.

[12] R. Bewley *et al.*, "Endangered archaeology in the Middle East and North Africa (EAMENA): Remote sensing and heritage management," *Antiquity*, vol. 89, no. 348, pp. 000–000, 2015.

[13] M. McClellan, "Automated classification of looting pits using satellite imagery," *Journal of Field Archaeology*, vol. 40, no. s1, pp. S89–S100, 2015.

[14] V. Prabhakar and R. Korisettar, "Threats to archaeological landscapes in peninsular India: A cultural resource management perspective," *Asian Perspectives*, vol. 53, no. 1, pp. 1–23, 2014.

[15] D. Stone, "Remote sensing and the documentation of archaeological site destruction," *World Archaeology*, vol. 42, no. 3, pp. 406–422, 2010.

[16] S. Parcak, *Satellite Remote Sensing for Archaeology*. New York: Routledge, 2009.

[17] C. Hammer and S. Ur, "Monitoring site damage in the Near East using multi-temporal satellite imagery," *Archaeological Science International*, vol. 10, pp. 1–12, 2014.

[18] N. S. Opitz and D. Cowley, "Advances in remote sensing approaches for archaeological prospection," *Advances in Archaeological Practice*, vol. 2, no. 4, pp. 000–000, 2014.

[19] Digital Archaeology Foundation, "Damage & Looting Detection Dataset for Afghanistan Archaeological Sites (DAFA–LS)," Zenodo, 2020. DOI: 10.5281/zenodo.3670055.

[20] K. Plessis, "Machine learning for landscape archaeology: A review," *Internet Archaeology*, vol. 52, pp. 1–20, 2019.